



Technical article

Risk is the world's leading financial risk management magazine. **Risk's** Cutting Edge articles are a showcase for the latest thinking and research into derivatives tools and techniques, with comprehensive coverage of the most recent advances in areas such as option pricing and hedging, market risk, credit risk, swaps and Monte Carlo methods.

www.risk.net/technical

Copyright statement

This PDF is made available for personal use only. Under the following terms you are granted permission to view online, print out, or store on your computer a single copy, provided this is intended for the benefit of the individual user. Should you wish to use multiple copies of this article, please contact reprints@riskwaters.com

© **Risk Waters Group Ltd**, 2003. All rights reserved. No parts of this article may be reproduced, stored in or introduced into any retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of the copyright owners.

Calculating the contribution

How do individual transactions contribute to the overall risk in a portfolio? Many models use the standard deviation of an economic factor, but one can also choose to work directly with marginal tail risk contributions. As Jack Praschnik, Gregory Hayt and Armand Principato demonstrate, such a choice has a considerable impact on the ranking of risk contributions

This paper describes a straightforward technique for allocating credit risk capital to individual transactions using Monte Carlo simulation. The principal advantage of the proposed technique is that it avoids reliance on standard deviation, while it retains the desirable property that risk allocations sum to total capital. It also demonstrates the importance of the method of choice for calculating risk contribution. For example, the amount of capital assigned to a transaction using portfolio standard deviation as the risk measure can differ significantly from the allocation derived from a conditional measure, such as that presented here.

Although our focus is credit risk, readers interested in a more theoretical treatment of the ideas presented here may refer to the articles on allocating value-at-risk (VAR) contained in the references.

Risk contribution is a term used to describe the risk of an individual asset within a portfolio. The idea is simply to take a portfolio-level measure of risk, which by definition incorporates covariance between exposures, and distribute it in such a way that the portfolio's risk is attributed to its constituent transactions. Because of the diversification achieved when an asset is combined with others, its risk contribution will be less than it would be when measured as a stand-alone asset.

The standard deviation measure

The primary use of credit risk contribution measures is the allocation of capital to individual transactions or business lines. When the focus is on capital allocation, the 'risk' being allocated is typically some high percentile of the loss distribution – for example, the 99th or 99.9th percentile – but one could allocate standard deviation or some other measure of risk.


In fact, there is no generally accepted definition of how to calculate risk contribution, although in practice most credit portfolio models base the calculation on the standard deviation of the portfolio.¹

There is a reason for the prevalence of the standard deviation-based approach. By construction, standard deviation can be allocated to individual transactions such that the sum of the individual risk contributions is equal to the total.

However, this 'summing' property does not hold if, for example, a transaction's risk contribution were to be defined as the change in the loss amount associated with, say, the 99th percentile of the loss distribution when that transaction is removed from the portfolio. Marginal capital, as this method is typically called, is a measure of the impact of an asset on the tail of the loss distribution, but marginal capital will not sum to total capital at a given percentile.

We have observed that most portfolio managers have a strong preference for risk contribution measures that have the summing-up property, which explains the widespread use of standard deviation. However, we also observe that portfolio managers prefer to allocate their economic capital – that is, risk measured at a high percentile – because this is the basis on which their performance is typically measured and their capital charges assessed. In order to meet both objectives, most credit models have adopted the approach of calculating risk contribution using standard deviation, which satisfies the summing criteria. That measure is then scaled up using a multiplier such that the individual risk contributions sum to economic capital. The multiplier is simply the ratio of economic capital to portfolio standard deviation.

Figure 1: Defining a neighbourhood around economic capital



This scaling approach makes a significant assumption. It assumes that the relative risk of the transactions in the portfolio is the same whether calculated in the centre of the distribution using standard deviation or at the tail of the distribution using economic capital. By relative risk, we mean that the ranking of transactions by risk is unchanged. However, the relative risk of transactions can change as one moves into the tail of the distribution, and therefore the amount of capital attributable to each transaction will also change.

Ideally, economic capital would be allocated using a direct measure of risk in the tail that maintained the summing property. While analytic solutions to this problem have been proposed for VAR, they necessarily make distributional assumptions that are less appealing for credit distributions. A simple and direct alternative is to employ Monte Carlo simulation. One way of doing so is described below.

An empirical allocation of economic capital

The direct allocation of economic capital to individual transactions requires an estimate of each transaction's conditional expected loss. The expectation is conditional because only large portfolio losses are to be considered in the expected loss calculation. Large portfolio losses are defined as those falling within a few percentiles of economic capital, as illustrated in figure 1.

In figure 1, the institution has defined economic capital at the 99.95th percentile, and, for the purposes of allocating this loss to individual transactions, will use information from all portfolio losses between the 98th and 100th percentile. A neighbourhood around economic capital, rather than a specific capital value, is necessary in order to gain enough observations from which to estimate tail allocations.²


¹ For example, standard deviation is the basis for credit risk contribution and capital allocation in models from KMV, RiskMetrics, and CreditRisk+, although it should be noted that RiskMetrics offers other choices as well.

² Note that one could place a tighter band around capital by choosing an upper bound of say, 99.99%. The results are not sensitive to the definition of the neighbourhood. However, a wider band results in more observations in the tail for a given number of iterations.

Table 1: Distribution of transactions and exposures in the test portfolio

Rating	Obligors	Average face value	Average exposure	Minimum exposure	Maximum exposure
Aaa	11	74,636,364	26,209,091	22,200,000	33,000,000
Aa	15	63,533,333	21,493,333	1,200,000	29,700,000
A	25	58,560,000	22,652,000	7,200,000	32,500,000
Baa	31	40,774,194	18,041,935	600,000	31,000,000
Ba	15	33,333,333	16,546,667	2,400,000	29,000,000
B	3	16,000,000	9,600,000	9,000,000	10,200,000
All	100	50,500,000	20,133,000	600,000	33,000,000

Figure 2: Tail risk contribution v. central (standard deviation-based) risk contribution for the entire portfolio



Calculation of the conditional loss requires that the simulation capture the number of times that each transaction defaults, conditional on the losses for the entire portfolio being in the neighbourhood of economic capital. The number of times a transaction defaults when total portfolio losses are in the defined neighbourhood, divided by the total number of portfolio loss events that belong to the neighbourhood, determines the conditional probability of default for each transaction. These conditional probabilities measure the probability that the transaction under consideration will contribute to a large portfolio loss.

The product of the conditional probability for exposure *i* and its estimated loss given default, LGD(*i*) gives a dollar risk contribution for the transaction based on the tail of the loss distribution:

$$\text{Tail risk contribution}(i) = \text{Conditional probability}(i) \times \text{LGD}(i)$$

The loss given default can also be made conditional on losses in the neighbourhood of economic capital. Here, we simply use the unconditional loss given default. Because the conditional probabilities are calculated over a neighbourhood in the tail region, the tail risk contributions will sum to a value close to, but not exactly equal to, economic capital.

However, expressing each obligor's tail risk contribution as a percentage of the total and multiplying by economic capital can achieve the desired summing property. This minor scaling does not change the relative risk contributions of the individual transactions, but it does result in a slight change in the dollar amount of capital allocated to each transaction.

Results for a simple portfolio

Tail risk contribution measures can be calculated for any credit portfolio

³ Each asset's default rate is described by placing a weight of 70% on the same systematic factor and 30% on an idiosyncratic factor.

model that employs a Monte Carlo simulation. To illustrate this, we consider a hypothetical 100-asset portfolio, the characteristics of which are summarised in table 1.

We examined the hypothetical portfolio using a simple credit portfolio model with the following characteristics:

- Transactions either default or not over the life of the credit exposure.
- The cumulative default rate for each transaction is log-normally distributed with mean, *P*, and volatility, *s*, derived from rating agency Moody's published corporate default data and the transaction's credit rating and assumed maturity.
- Loss given default is assumed to be constant, for reasons of simplicity.
- Transaction default rates are correlated via a shared relationship to a single systematic risk factor.

To tabulate the credit loss distribution, we used a standard Monte Carlo simulation, where the default probability for each transaction on a given iteration was determined by a draw from a multivariate log-normal distribution. If we use the 99.95th percentile to calculate economic capital, our hypothetical portfolio would require economic capital of \$184,600,000 (9.2% of the exposure).

For each obligor, we calculated the tail risk contribution measure, using the 98th to 100th percentile neighbourhood. We also calculated two standard deviation-based risk contribution measures for comparison:

- Undiversified risk* – the standard deviation of the transaction's expected loss without taking into account covariance with other assets. Assuming that defaults are binomially distributed,

$$\text{Undiversified risk}(i) = \text{Sqrt}(P \times (1-P)) \times \text{LGD}(i)$$

- Central (standard deviation-based) risk contribution* – standard deviation is calculated using the obligor-specific default rate volatility and the assumption – used in the simulation – that all assets share the same relationship with the common systematic factor.³ This measure is then scaled up to economic capital following the practice commonly employed in commercial credit models – that is, the amount of capital is increased so that economic capital is fully accounted for, but the relative risk contributions remain unchanged.

All three risk contribution measures are summarised in table 2, where we set out the bottom, middle, and top 10 transactions as ranked by their tail-based risk contributions. (Table 2 also provides the exposure, expected default rate, and default rate volatility of these transactions.)

The relative risk rankings for the entire portfolio are plotted for both the tail-based allocations and the central (portfolio standard deviation) allocations in figure 2 (1= lowest risk, 100 = highest risk). Taken as a whole, the three different risk contribution calculations are fairly similar. For example, all three identify, to within a few ranks, the 10 highest and lowest risk contribution exposures. The rank correlations calculated over all 100 transactions exceed 0.95 for all three pairings of risk measures.

Across all three risk measures we assume the following: transactions with a low mean default rate and low volatility will have low conditional probabilities of default when portfolio losses are large. Transactions that are highly correlated with other transactions in the portfolio will have higher conditional probabilities and therefore higher tail-based capital allocations.

Nonetheless, within the 100 transactions, individual ranks and therefore individual dollar risk contributions differ substantially, particularly in the middle of the distribution. This differential is the direct result of measuring the tail risk rather than the central, standard deviation-based, risk.

The dollar risk contributions from each method can be compared after normalising all three methods to total capital of \$184.6 million.⁵ The average, absolute deviation between the tail-based measures and portfolio standard deviation-based measures is \$449,704 or about 24% of the average risk contribution of \$1.846 million. These are significant differences when allocated capital is used to calculate risk-adjusted return or for internal capital charges.

Conclusion

One approach to the calculation of risk contributions and capital allocation for a credit portfolio involves allocating portfolio standard deviation to individual transactions – or exposures – and scaling this measure up to obtain an allocation of economic capital.

The standard deviation approach is reasonable in many situations, and straightforward to calculate. However, it implicitly assumes that an asset's contribution to standard deviation is proportional to its contribution measured at the extreme tail of the distribution – a less plausible assumption for the allocation of credit risk.

Alternatively, capturing results from the Monte Carlo simulation and empirically estimating the probability that each transaction contributes to large, capital-level losses could serve to estimate risk contributions. The methodology described in this article is one way of implementing this approach. The tail-based risk contribution measure highlights the fact that risk contribution is a function of the confidence level chosen for risk.

For higher-percentile capital estimates, it becomes increasingly likely that the lowest-risk deals will have defaulted in order to have resulted in such a large loss. As one moves toward higher confidence levels for economic capital, the risk contributions of the lowest-risk exposures increase.

Because relative risk contributions change with the percentile used for

⁴ Although risk contributions calculated via simulation can sometimes be volatile, we believe the differences in table 2 are not due simply to volatility in the estimates. Because the portfolio is small and the number of iterations large, the risk contributions based on both standard deviation and tail based estimates are stable.

⁵ The 99.95 percentile loss our hypothetical portfolio.

economic capital, they also change relative to central, standard deviation-based measures.

Neither of the approaches is appropriate for all applications. Rather, portfolio managers should be careful to define what they want to accomplish with risk contribution measures, and should be aware that the methodology for calculating them can make a significant difference to the final allocations. We believe the tail-based measure is the most appropriate risk allocation measure for determining how much economic capital to associate with a specific transaction. □

Jack Praschnik, Rutter Associates
Gregory Hayt, Rutter Associates
Armand Principato, MBIA

Bibliography:

A new approach to component Var, Carroll, RB, Perry, T, Yank, Z, Ho, A, Journal of Risk, Volume 3 No 3, Spring 2001

CreditMetrics Technical Document, Gupton, G, Finger, C and Bhatia, M, RiskMetrics Group, 1997

Decomposing Portfolio Value-at-Risk: A General Analysis, Hallerback, W, Erasmus University, Working Paper, December 29, 1999

Table 2: Risk contribution and relative risk ranking for the top, middle and bottom portfolio tiers

Obligor	Credit quality	Exposure (US\$)	Average life	Default probability*	Default rate volatility	Method for allocating economic capital to individual transactions					
						Tail risk contribution		Undiversified risk		Central risk contribution (standard deviation)	
						Percentage**	Rank†	Percentage	Rank	Percentage	Rank
2	Aa	1,200,000	2	0.03%	0.11%	0.0002%	1	0.0089%	5	0.0033%	5
81	Aaa	24,300,000	1	0.00%	0.00%	0.0004%	2	0.0000%	1	0.0000%	1
82	Aaa	24,600,000	2	0.00%	0.00%	0.0004%	3	0.0000%	2	0.0000%	2
83	Aaa	24,900,000	3	0.00%	0.00%	0.0004%	4	0.0000%	3	0.0000%	3
53	Aaa	26,500,000	3	0.00%	0.00%	0.0004%	5	0.0000%	4	0.0000%	4
1	Baa	600,000	1	0.14%	0.28%	0.0011%	6	0.0091%	6	0.0041%	6
3	Aa	1,800,000	3	0.07%	0.22%	0.0011%	7	0.0201%	7	0.0095%	7
41	Aa	20,500,000	1	0.02%	0.11%	0.0034%	8	0.1205%	9	0.0556%	9
51	A	25,500,000	1	0.02%	0.11%	0.0042%	9	0.1499%	11	0.0691%	13
12	A	7,200,000	2	0.06%	0.11%	0.0064%	10	0.0722%	8	0.0196%	8
8	Baa	4,800,000	8	3.68%	2.11%	0.28%	46	0.37%	27	0.25%	32
37	A	18,500,000	7	0.90%	1.10%	0.28%	47	0.72%	49	0.50%	55
85	A	25,500,000	5	0.50%	0.73%	0.29%	48	0.74%	51	0.45%	53
24	Baa	14,400,000	4	1.39%	1.09%	0.29%	49	0.69%	47	0.38%	44
86	A	25,800,000	6	0.69%	0.90%	0.29%	50	0.88%	59	0.56%	59
47	Aa	23,500,000	7	0.65%	0.74%	0.31%	51	0.77%	54	0.42%	49
62	Baa	31,000,000	2	0.44%	0.42%	0.32%	52	0.84%	58	0.32%	39
34	Baa	17,000,000	4	1.39%	1.09%	0.34%	53	0.82%	55	0.45%	51
65	A	32,500,000	5	0.50%	0.73%	0.36%	54	0.94%	60	0.58%	61
68	Aa	20,400,000	8	0.85%	0.90%	0.38%	55	0.77%	53	0.45%	50
18	Ba	10,800,000	8	17.37%	10.68%	2.98%	91	1.68%	84	2.83%	89
19	Ba	11,400,000	9	19.71%	10.82%	3.01%	92	1.87%	87	3.03%	90
17	B	10,200,000	7	31.08%	15.02%	3.21%	93	1.94%	89	3.77%	94
20	Ba	12,000,000	10	21.55%	10.94%	3.75%	94	2.03%	90	3.23%	91
55	Ba	27,500,000	5	10.40%	7.64%	4.26%	95	3.45%	96	5.20%	95
69	Ba	20,700,000	9	19.71%	10.82%	5.47%	96	3.39%	95	5.55%	96
50	Ba	25,000,000	10	21.55%	10.94%	7.82%	97	4.23%	99	6.82%	98
56	Ba	28,000,000	6	12.85%	8.91%	7.92%	98	3.85%	97	6.20%	97
58	Ba	29,000,000	8	17.37%	10.68%	8.00%	99	4.52%	100	7.75%	100
57	Ba	28,500,000	7	14.99%	9.91%	9.18%	100	4.18%	98	7.05%	99

* Default probabilities and volatilities are derived from Moody's historical data.

** Risk contributions for all three methods as a percent of total economic capital (\$184,600,000)

† Higher ranks imply higher risk